XaCCT 3.0

Technology Overview

Revision 2.1

20/09/97

Confidential Information

Contents

Table of contents

Table of figures

Introduction

Manifest

XaCCT Technologies develops and markets the XaCCT product family for auditing, accounting, billing and planning of TCP/IP networks, targeted at handling the economics of TCP/IP networks on the Internet and corporate Intranets.

The company's revolutionary products analyze network traffic by gathering data from a wide range of network elements e.g. routers, hubs, firewalls, application servers, captured network packets and directory name servers. The data is then efficiently processed and stored using highly interactive configuration and reporting tools.

XaCCT provides the tools to generate a report similar to a telephone report by treating Network Sessions as telephone calls. A Network Session is a set of parameters that describes a complete client-server application session. These parameters may include IP addresses, user names, profit center names, application type, duration, information about the content of the session, etc.

The Need for Network Accounting and Billing

Low costs of Internet connectivity and a wide range of services are driving more and more people into the information highway and forcing companies to deploy TCP/IP networks. This process has led to a new market of client-server applications that enables the user to interact with other users and computer systems around the Net. These applications are gradually consuming Intranet and Internet bandwidth whilst computer networks are reaching their physical speed limits.

Organizations suffering most from uncontrolled network usage growth are the Internet service providers, corporations and other institutions using Wide-Area Networks (WAN).

Internet Service Providers need a profitable businesses model

E-mail, the flagship service of today's Internet, consumes very little network traffic, working efficiently even at offline environments or low priority bandwidth.

Classical client-server applications such as WWW and FTP are designed to consume bursts and peaks of bandwidth which may further be controlled by specialized hardware and software that serve to cache and regulate the flow of data requests.

Applications such as voice and video conference require the full capacity of Internet and Intranet bandwidth at all times. Whether the bandwidth comes from an Internet Service Provider's (ISP) dial-up connection or from a corporate Internet connection, it will be occupied to its capacity by these applications.

At present most ISP customers pay either a flat fee or are billed in relation to the amount of time spent on-line. Flat fee may be adequate for broadcast services as expenses do not vary according to the amount of usage. Internet users exhibit many different patterns of usage, from low consuming E-mail and web users to high consuming digital video and phone users. Maintaining the high bandwidth WAN is not financed in direct relation to users' network consumption and can therefore not be covered by a flat fee.

A European ISP may pay around $5000 per month for a leased line of 64Kbps connected to a US based Internet provider. This line renders an average effective throughput of about 6Kbytes/second. If this figure is multiplied by the number of seconds in a month we get the monthly throughput of 1.59e10 bytes. This gives the cost of about $0.33/Megabyte. A dial-up user using a voice chat or video conference software consumes an average of 2Kbytes/second using a 28Kbps modem which totals approximately 7Mbytes/hour. It therefore follows that a bandwidth consuming user can cost the provider up to $2.3 per hour.

By offering users "unlimited" usage or a fixed price of $40 for 30 hours (for example), the cost of the international line is not being covered by the average high-bandwidth user. The international line becomes over-occupied and customers approach competitors for better Internet connection throughput. Competition requires the ISP to upgrade the bandwidth of the Internet line thus reducing profitability margins.

XaCCT provides tools to bill customers in diverse ways according to the actual usage of network resources. This can only be done by accumulating information on Network Sessions from various parts of the WAN and linking them to users. Billing schemes can be configured to various models based on flat charges, on-line time, time of day, network load, resource utilization, bandwidth consumption, etc.

Corporate and Academic WANs need a cost distribution control system

Corporations and academic WANs are composed of network elements such as copper lines, fiber, HUBS, switches, routers, terminal-servers, modems etc. which are spread throughout departments, buildings, cities and countries, interlinked by Telcos and centrally maintained by skilled professionals. The cost of maintaining and upgrading these network elements may vary from tens of thousands to millions of $US per year. In most cases corporations share these costs and charge departments, users or any other profit and loss entities equally.

At present, corporate Internet connections are unaccounted expenses which results in employees being able to use the line for any purpose. Unlike a telephone billing report, there is no tracking mechanism for corporate bandwidth usage.

XaCCT provides the tools to bill profit and loss entities in diverse ways by accumulating information about Network Sessions from all parts of the WAN and then linking them with profit and loss entities. Spreading the costs is done on a per-usage basis rather than on a flat-rate basis.

Network billing vs. Telephone billing

The Telco market has proven to be one of the most profitable and stable markets during the last half century. This is largely due to the ability to produce comprehensive and detailed monthly customer billing reports.

Telephone bills usually contain the following details:

- The initiator's telephone number
- The destinations telephone number
- The time that the conversation started and ended

The tariff for the telephone conversations varies depending on:

- The time of day
- The route of the call (prefix, long-distance etc.)
- The type of call (toll free, etc.)
- The length of the conversation

XaCCT provides a framework for billing Network Sessions, analogous to telephone billing.

Network billing reports usually contain the following basic details:

- The initiator's IP address and associated user or entity name
- The destination IP address and associated user or entity name
- Type of Network Application
- The time that the conversation started and ended
- The amount of data transferred back and forth

The tariff for conversations varies depending on :

- The time of day or overall network load at the time of conversation
- The route of the transaction (LAN, WAN, Internet)
- The priority of the transaction (where applicable)
- The type of Application (Web, FTP, VDO, SQLnet, NFS, E-Mail, etc.)
- The length of the conversation
- The amount of data transferred

One of the major differences between network billing as compared to telephone billing, is that in network billing it makes much more sense to charge for bandwidth, type of service and priority in addition to time.

XaCCT 3.0 Basic Features

XaCCT 3.0 provides a unique view of network activity throughout the WAN. This view is based on records of traffic activity called "Network Sessions" that are analogous to telephone calls records (Call Data Records (CDRs))

The system is composed of two parts: a) a distributed, highly scaleable data collection and database repository system, and b) applications that processes the Network Sessions gathered in the database. Both parts of the system are configured through a user-friendly web-based interface.

XaCCT 3.0 produces comprehensive network accounting and billing reports for internal entities such as users, departments and any profit and loss entities. The reports can also include information about external entities such as a company name, a contact person, an address or category. Internet distance may be determined for remote Internet sites which have exchanged data with local hosts. The information is collected by cross-referencing various sources of network related information which may be both internal and external to the organization.

XaCCT can be used for the following purposes:

- Implementation of new economic models and innovative billing schemes for network utilization and Inter-Network communications for different entities.
- Accounting and billing of sub-organizations for their specific use of organizational networking resources.
- Accounting and billing for network traffic produced by users and organizations connected to an ISP.
- Tracking of possible misuse (or abuse) of expensive corporate networking resources by employees.
- Logging of network usage for the purpose of back-tracking network events whilst trying to analyze network failures or breaches of security ("postmortem" intrusion analysis).
- Analyzing the use of corporate networking resources in order to enable effective base-lining, back-charging and cost allocation for the upgrading and scaling of these resources.
- Predicting trends in network consumption and simulating growth in network usage due to introduction of new bandwidth-hungry applications such as CTI (Computer Telephony Integrated).
- Invocation of alerts upon the occurrence of critical network related events.
- Analysis of network traffic based on the content level. This type of application may be useful for various business applications like help-desk, customer services, call centers and web-based commerce.
- Integration of a new network session data collection technology with legacy accounting, billing and customer support systems in both the corporate and the Telco environments.

System Overview

XaCCT 3.0 consists of a distributed data collection system and several applications that utilize the collected data. The system is typically installed once on several computers throughout the WAN and extended to other computers as the WAN grows. Upgrading and configuration is centrally administered without need for human intervention at remote computers.

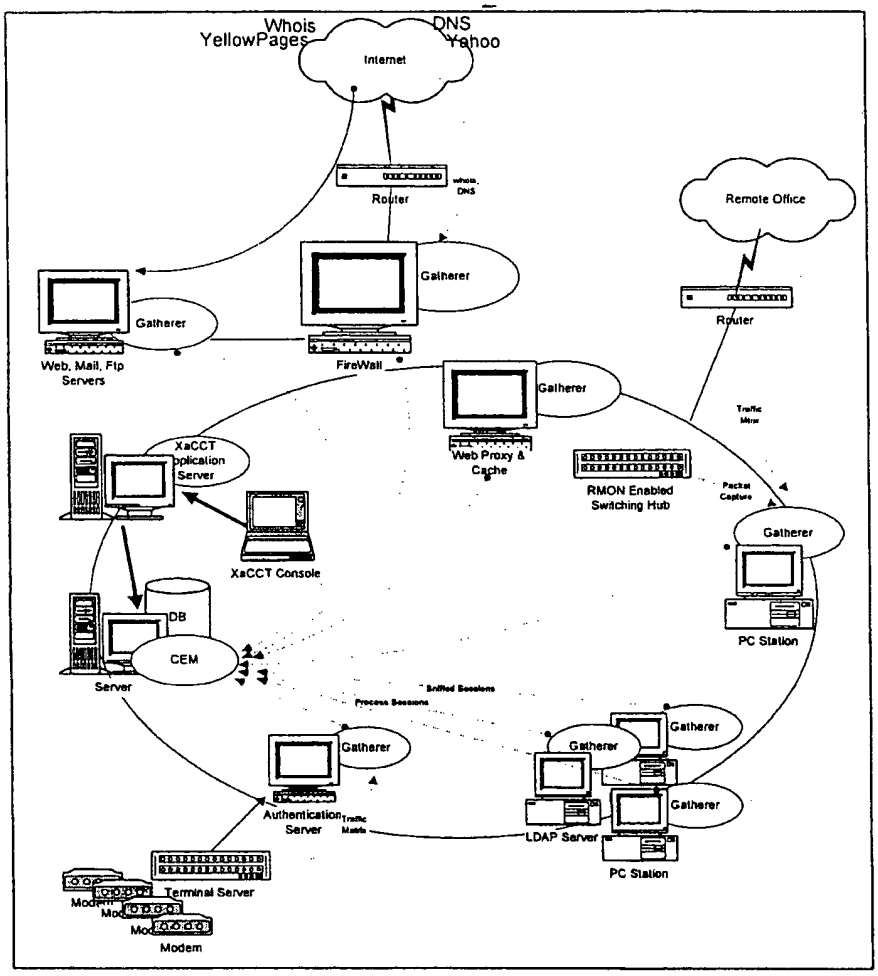

Figure 1 - Schematic overview of XaCCT

Gatherers probe an abundance of information sources extracting Network Sessions, real-world names name services and other databases. They are strategically located throughout the WAN close to information sources which they probe. Information collected is compressed and transformed into Unified Network Information Records (UNIRs). The distributed pieces of information consolidated by means of an enhancement process associating the UNIRs to world entities.

The CEM (Central Event Manager) receives enhanced UNIRs from the distributed Gatherers. It controls the enhancement process, merges various sources of UNIRs in order to reduce redundancy and duplications and stores them in a database. The whole data collection system is configured using a web-based application.

The XaCCT applications utilizes the UNIRs collected in the database to provide various solutions. The user uses a Java enabled web browser to utilize the application. The applications are used to perform various auditing, tracking, planning, billing and accounting tasks.

Gathering Information about Network Sessions

Introduction

XaCCT 3.0 collects information from a large variety of information sources in order to provide a complete view of the activity throughout the WAN. These information sources are either involved in transporting Network Sessions or in providing information about the parameters of the Network Session.

Example1 : Consolidating Multiple Information Sources in a Corporate WAN

The following is a list of network information sources from which information can be collected by XaCCT 3.0 in order to produce network billing and accounting reports:

- A Firewall configured to log all Network Sessions that pass through it.
- Web proxy servers that log their traffic.
- The internal corporate DNS (Domain Name Server) containing name identifications of participants in a Network Session.
- The corporate E-mail server logging information about every E-mail sent and received whether it is located within the corporate network, or channeled to/from the Internet.
- Directory services associating E-mail addresses to individual users and departments.
- Routers and switching-hubs may provide network consumption statistics.

By consolidating the information from all these sources it is possible to associate correctly between real-world entities (e.g., user names, departments) and network activity.

Using the XaCCT Reports User Interface, this information is transformed to a billing report which charges departments and individuals for web access through the proxy, direct Internet access, inter-departmental traffic, time of day, "expensive" protocols etc.

Example 2: Consolidating Multiple Information Sources in an ISP WAN

In addition to the above list of network information sources, an ISP makes extensive use of the following information sources. By consolidating information from all these sources, an ISP can produce effective billing reports.

- TACACS or RADIUS authentication servers provide information on individual users logged in via modems or ISDN connection.
- Terminal-Servers may provide network consumption statistics.

The ISP's network information sources together provide a complete picture of the activity of users consuming network resources.

Using the XaCCT Reports User Interface information is converted into a billing report which charges individuals by duration, time of day, network load, web access through the proxy, direct Internet access, ISP customer to customer communication, "expensive" protocols etc.

Information Source types

Most network devices and application servers provide logging or statistical information about their activity. A network information source can be the log file of a mail server, the logging facility of a Firewall, a traffic Statistics table available on a Router and accessible through SNMP, a database entry accessible through the Web, an authentication server's query interface etc.

Depending on the type of information or log, a Gatherer contacts the information source by whatever means necessary.

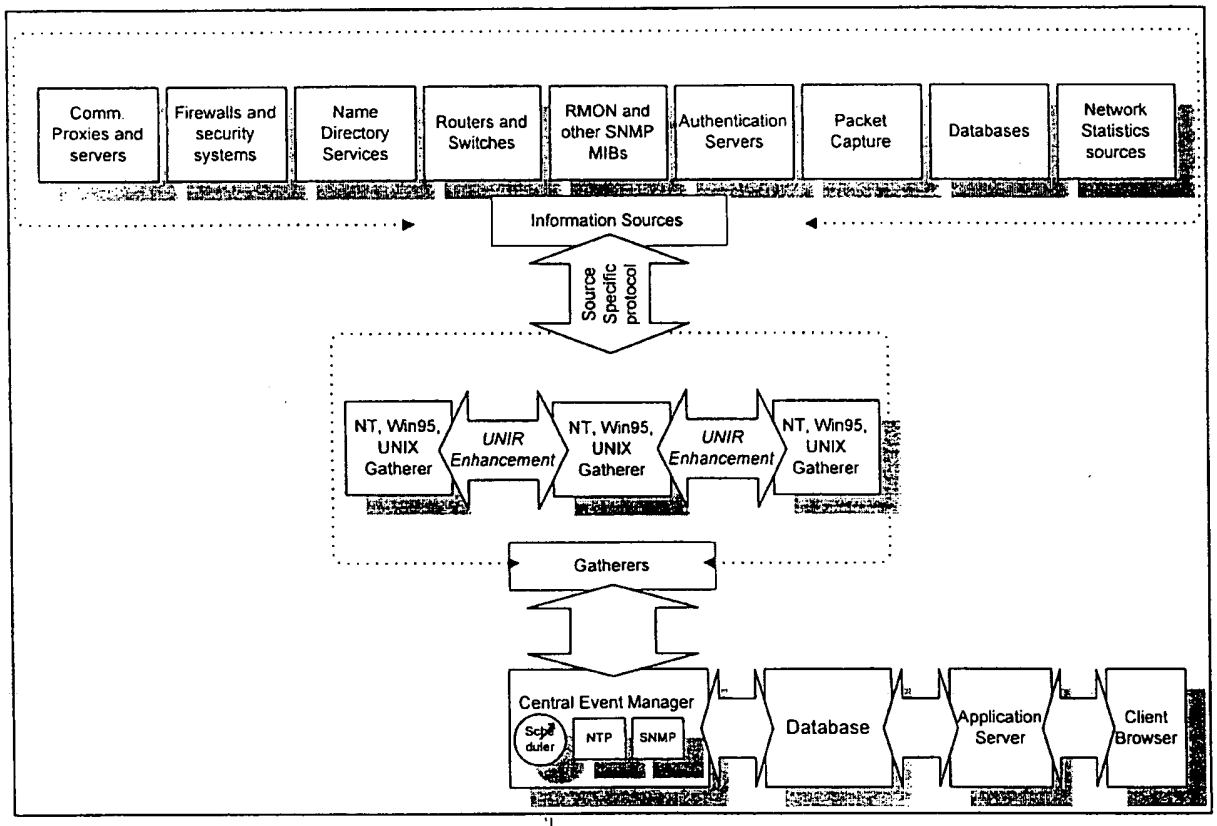

Figure 2 - multi-type information sources

XaCCT treats Information Sources as either synchronous or asynchronous. A synchronous I.S. serves as a trigger to the XaCCT collection system. A UNIR enhancement process (which merges parameters from several IS's) begins with an asynchronous I.S. The records from an asynchronous I.S. are transferred to another Gatherer which probes this time a synchronous I.S. in order to add the missing parameters. This enhancement procedure can repeat several times until all the wanted records have been collected and the UNIR is placed in the database.

There are exceptions where an asynchronous I.S. is used as synchronous source. For example, a Radius server continuously delivers users login/logout records. If the system were to collect CISCO NetFlow records and to enhance the records adding user names to the IP address found, using data from the Radius server, the XaCCT collection system would take the NetFlow Information Source Module (in a Gatherer) as an asynchronous information source and the Radius Information Source Module (ISM)as a synchronous I.S. In this case the Radius ISM will have to continously receive Radius records and "pretend" to be a synchronous information source by delivering the user name associated to an IP address, upon demand.

Network Statistics Information Sources

Some applications and network devices provide statistical information about network traffic flow. By accessing these information sources regularly, it is possible to generate informative Network Sessions UNIRs.

Examples of Network Statistics sources:

- CISCO IOS 10.X accounting - available in every Cisco Router
- Bay Networks RMON
- CISCO RMON
- Application RMON Probes
- NetRaMet (DOS, UNIX)
- NetFlow (CISCO)

Network Sessions Information Sources

The most informative UNIRs are produced from access logs of server-applications and network-activity logs of some firewalls. Application servers know all about the Network Session in which they are involved. Firewalls using advanced packet filtering technology know the details of any Network Session that passes through them in order to let through packets that are part of the transaction and to avoid letting through any unrelated packets. Firewalls using proxy technology know all the details of a Network Session that is performed through the proxy.

Examples of Network Session sources:

- FireWall-1 2.X (UNIX,NT) + XaCCT 2.0
- FireWall-1 3.X (UNIX,NT)

- FTP server (UNIX)
- INN USENET server (UNIX)
- PIX Firewall (CISCO)
- Alta Vista Firewall
- Raptor Firewall
- Sendmail (UNIX)
- Smap (UNIX)
- WWW IIS Access (NT)
- WWW Netscape server and proxy (UNIX/NT)
- WWW Cern, Apache server and proxy (UNIX)

Packet Capture Information Sources

Information sources that probe the network and capture every packet flowing, provide the basis for generating high quality UNIRs.

The XaCCT 3.0 Packet Capture Module can deal with various types of complex protocols sometimes requiring a look into the content of the packet in order to be able to associate packets with different port numbers of the same Network Session. FTP uses a different port number for every file transferred during an FTP session. The port number may be found as text in the data of the first packets of a file transfer.

XaCCT 3.0 UNIX/Win95/NT Packet Capture Module is one of the possible packet capture sources. However, an RMON probe inside a network device such as a switching hub or a router, "sees" the packets that flow through the device. An RMON probe therefore provides information about several network segments simultaneously, while a packet capture application may only see one segment at a time.

Examples of available packet capture sources:

- Bay Networks RMON
- 3COM Switching Hub RMON
- CISCO 7000 RMON
- SunScreen Packet Vectoring
- Network General Sniffer®
- Frontier Software RMON Probes

User and Host Identification Information Sources

There are different technologies for identifying and authenticating users and computers. User authentication systems require the user to be identified before accessing a network resource. In some cases it is possible to associate a user with an IP address without the user being authenticated, for example when a corporate user is

the sole user of a PC. User identification implies the association of an IP address to a user or any other real-word entity.

Examples of available auxiliary information sources:

- RADIUS (Shiva, etc.) - User authentication used with dial-up connections.
- TACACS(+) - User authentication used with dial-up connections.
- Ident - User identification standard protocol for an active TCP/IP connection.
- pcnfsd - User identification server used with Sun PC/NFS software.
- Kerberos - User identification standard used by some UNIX systems.
- DNS - Internet Name Server, maintains IP - Name databases.
- WINS/DHCP - Microsoft IP allocation and IP - Name database.
- WHOIS - Name/Company Lookup databases accessible through the Internet.
- SecuRemote (CheckPoint) - Remote access User Authentication.
- SKIP - Authentication standard.
- CryptoCard - Remote access User Authentication.
- SecureID - Remote access User Authentication.
- S/Key - Authentication standard.

Auxiliary Information Sources

XaCCT 3.0 provides intelligent methods of analyzing information about remote Internet sites which have exchanged data with local computers. By cross-referencing information sources, reports are produced where remote sites are listed not merely as bogus names or IP addresses, but as company names, contact people, addresses, categories and network distances.

Examples of available auxiliary information sources:

- Yahoo Category Index - to find category of an Internet domain name.
- Dun & Bradstreet - to find parameters in the profile of a company
- Webster Control List™ - to find category of an Internet domain name.
- DNS - associates a domain name to IP addresses and vice-versa.
- WHOIS - find the name of a company or person.
- traceroute, AS-traceroute - find the Internet distance and effective throughput

Technology

Figure 3 - Collection System Top-level Design

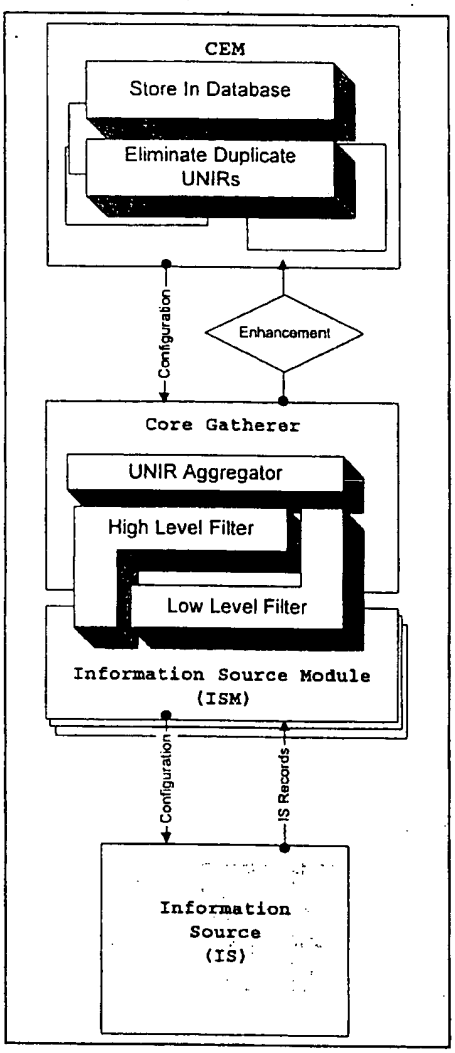

The Gatherer

XaCCT obtains Network Sessions and naming information from distributed "smart agents" called Gatherers that probe the information source and generate UNIRs (Unified Network Information Records that contain a compact representation of the total session content).

Scaleable and Light-Weight

Gatherers are multi-threaded and lightweight, designed to run on non-dedicated hosts as background processes. Gatherers are located close to the information source they access in order to reduce traffic across the WAN. Each Gatherer can collect information from multiple sources, allowing scaleable configuration of the collection system to accommodate for large enterprise networks.

Supports Multiple Information Source Types

Gatherers may access information sources using several methods including SNMP, TCP & UDP client/server, File, Telnet, ODBC, DNS, command execution and FTP. This allows for connections to a wide array of existing and future information sources.

Centrally Administered and Upgraded

Gatherers are administered by the Central Event Manager (CEM), allowing for central administration of the entire distributed XaCCT system, reducing maintenance cost and complexity. The Gatherer software itself, after initial installation, is automatically upgraded in the field, without need for any manual administration.

Supports Data Caching

The Gatherers utilize a local cache of information obtained from information sources in order to minimize the access to resources such as CPU and network. This is a learning process that is designed to minimize network consumption and to speed up work.

Extensible

XaCCT is designed to integrate with various types of 3rd part products that emit network related information. It implements a unique technology called *network session enhancement* where information is channeled between gatherers until reaching the database. With each stop in the channel more information is accumulated about the network session. New or updated information sources can easily integrate with the XaCCT system. cartridge-like software modules use a unique API (Application Programming Interface) to integrate with the collection system. Each new information source is configured and probed by an ISM (Information Source Module) which implements the API. (Each ISM comes with an ISCM - Information Source Configuration Module, which provides the user interface for configuring the Information Source). In the future, the API will be exposed to the public and enable customers and OEMs to add proprietary ISMs and their adjacent ISCMs and extend XaCCT to meet various market segment needs.

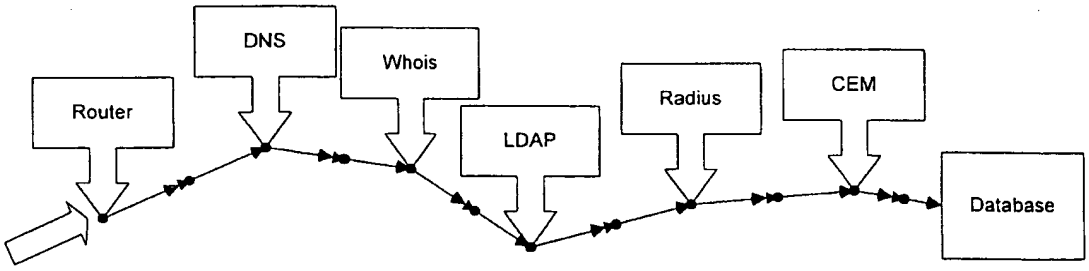

Figure 4 - Enhancement Procedure

Fault tolerant

The Gatherer can handle critical situations such as loss of connection and restarts by storing data on disk using a local backup mechanism until things are back to normal. After a given period of time, the collection process stops if necessary and reset when the communication is back. Information Source Modules are automatically cached on disk.

Passes through FireWalls

Gatherers can communicate through proxy gateways, firewalls and address translation barriers by implementing an internal routing mechanism.

Communicates Securely

Communication between Gatherers and communication between Gatherers and the CEM is authenticated and encrypted using industry standard 3rd party solutions, allowing sensitive information to be passed without jeopardizing overall system security.

Does not consume much bandwidth

The amount of network traffic generated between Gatherers and with the CEM is configurable by the user. This will depend on the level of detail and level of aggregation needed and determined by the user, and by the types of available information sources. The user may configure filters to avoid tracking of unwanted Network Sessions. The user may also define the granularity of details and aggregate sessions with specific protocols, users, networks etc. A typical consumption of 1-5% should be exhibited in LANs that hold a Gatherer.

Figure 5 - Filtering and Aggregating

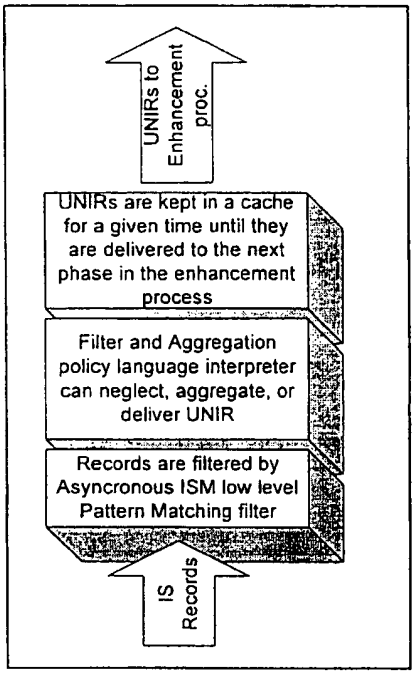

In order to reduce XaCCT generated network traffic and lower the quantity of data stored in the database, XaCCT Gatherers facilitate to filter and aggregate UNIRs. The filtering mechanism neglects UNIRs that are not needed for in the collection process. This is done by providing the Gatherer with a filter expression. The expression comprises of a Boolean expression containing patterns to be matched against parameters that are collected from a given Information Source. This expression is the *Filtering and Aggregation Policy* and can be configured by the user of the system for any phase of the UNIR enhancement process.

The UNIR Filtering and Aggregation Policy language enables the XaCCT application user to configure policies that will match some of the UNIRs, will then either disregard them, aggregate them or send them on as they are.

The aggregation mechanism takes several UNIRs that match a given expression and transforms them into a single UNIR, aggregating common parameters during this process, but possibly loosing some other fine-grained details. For example, when downloading a web page, separate UNIRs are generated for every piece of graphics, text and program that appear on the page. In most cases, this level of detail and granularity of information is not needed for the purpose of accounting and billing. It would suffice to have one UNIR representing the entire session of downloading the page and a record stating the number of UNIRs that were aggregated in the process, the total number of bytes transferred etc.

The Central Event Manager (CEM)

The Central Event Manager (CEM) orchestrates the whole collection, processing, and storing process of the XaCCT 3.0 system. Data collected and processed by the Gatherers is stored by the CEM in an external database.

The CEM should be run on computer with sufficient memory, CPU and networking capabilities to support the particular system installation.

Eliminates Record Duplications

The CEM is responsible for consolidating information from multiple sources and eliminating accounting duplications. If a transaction is initiated between office in a WAN and another office, the session may be logged twice by Gatherers on each side. The user can configure the conditions for merging the UNIRs from both Gatherers and the consolidation function. Transaction consolidation may also serve to further enhance the information about a transaction.

Performs Regular Data Aging Procedures In the Database

Once the transaction is stored in the database it stays there until the CEM expires its validity due to time or space limitations.

Initiates and Monitors System Performance

The CEM allows for central configuration and administration of the entire system of Gatherers and information sources. The CEM continuously monitors the state of each Gatherer and information source in the system, alerts when a Gatherer is unreachable and resets the collection process when a Gatherer is reachable again.

Centralized system-wide Upgrade, Licensing and Data Security

Upgrading the system is performed only once, at the CEM. Gatherers are automatically upgraded once a new version is installed with the CEM. Information Source Modules (cartridges) are also exported via the CEM to the Gatherers. The CEM is also responsible for implementing the licensing policy and also deciphers encrypted data from remote Gatherers.

Configurable Through a Web-Based User Interface

The CEM is accessible through a Java enabled Web browser. The configuration is password protected, allowing only authorized personnel access and configuration and sensitive information.

The system in configured using a set of wizards, guiding the user step-by-step through the installation and configuration processes. On-line help is provided at every step for additional guidance.

The user can monitor an information source in order to view its current status and activity log.

The Unified Network Information Record (UNIR)

The UNIR is a set of parameters related to network sessions. These parameters can be used for various accounting, auditing, billing, security and network management applications.

In an ideal world, it would be possible to obtain the following parameters about every network session through standard APIs supported by network devices and application servers:

Parameters	Description	Category
SrcIP, SrcInternal, SrcXlated, DstIP, DstInternal, DstXlated	Source IP Address (session initiator) Is the Src internal ? Has the Src been addr. Translated ? Destination IP Address Is the Dst internal ? Is the Dst translated	*IP related information*. Information associated to the addresses of the source and destination of the session. Also associated with addresses are flags indicating weather the address is internal or external to the network and weather it has undergone address translation.
SrcResolved, DstResolved, SrcIdent, DstIdent, SrcCategory, DstCategory ⇒ Consumer	DNS resolved Source IP DNS resolved Destination IP Identification assoc. to Srouce IP Identification assoc. to Dest. IP Category associated to Source IP Category associated to Dest. IP Who gets billed for this session	*Identity of users and entities*. Information associated to the source and destination addresses, such as identity of the person or asset and category of an external entity. The Consumer may be any unique identity associated to the transaction and must be identified in the billing system.
Service, SrcPort, DstPort, Protocol	Application name Source port number(s) Destination port number(s) TCP, UDP, etc.	*Service or application associated to session*. A single application can generate multiple network flows with different source and destination ports and protocols (TCP, UDP etc.)
StartTime, EndTime, TimeOut	Start of session End of session Did the session end properly ?	*Time*. The start and end times of a session and weather or not the session ended properly.
InPkt, InBytes, OutPkt, OutBytes, TotPkt, TotBytes	Incoming packets Incoming bytes Outgoing Packets Outgoing bytes Total number of packets Total number of bytes	*Data transferred*. Information about the amount of total and net traffic generated during the session.
Throughput, Retrans, Load, Latency, Priority, RSVP Hops, AsnHops, Aggregation	Effective bytes/sec for this session Signal to noise ratio Network load during session Lag b/w client and server packets Router priority for this session RSVP applicable for this session Distance in hops (to destination) Distance in Autonomous systems Number of UNIRs compressed	*Quality of service*. Parameters associated to quality of service such as effective throughput, retransmission statistics, network load at the time, router priority of the session and network geography parameters. If several UNIRs were compressed into a single UNIR in order to reduce network volume, the number of UNIRs will be in the Aggregation param.

Price, PriceCategory	Real-time calculation of price Price category associated to session	*Accounting parameters.* Parameters related to billing and accounting, are sometimes calculated in real-time rather then at a post-processing stage.
Content	Depends on the type of session.	Service-specific data about the content of the session such as MIME attachment names for e-mails, file names for FTP and NFS, URLs for HTTP etc. Future versions will permit recording of session content. Recording session content will also permit filtering and aggregating based on this information.
Gatherers, ISMs, InfoSrc, Hosts, Locations	Gatherers during enhancement process Information Source modules Information Sources Hosts Physical locations	*Collection process parameters*. The UNIR enhancement process involves several stages on possibly different hosts, these parameters include the Gatherer name, host and location, the Information Source and program module that handled it. These parameters may contain multiple values.

Unfortunately, in today's networking environments, it is not always possible to obtain all these parameters for a given network session. In addition, parameters related a given session may be dispersed over many network components. To obtain these pieces of information, it is required to access many different information sources for every single session.

XaCCT 3.0 collection system mediates for this lack of industry standards for network accounting, by accessing any logging, tracking, accounting and statistical information sources and transforming these pieces of distributed information parts into a unified record format called UNIR.

UNIRs and Network Flows

A solution to the problem of providing an in-depth knowledge of what's going on in the network and representing it in compressed informative records, has been proposed and implemented in the last few years by an academic work-group lead by Nevil Brownlee from the University of Auckland in New Zealand [1,2] and lately also implemented by CISCO Systems [3].

The idea is to count the number of packets and number of bytes associated to packets with matching parameters (such as address and ports) and to export these network "flows" periodically to a collection system. Network flow "meters" run on computers located on LANs where traffic is being measured, or inside network devices such as routers.

Network Flows collection will be supported by XaCCT in early releases (aswell as Cisco FlowCollector's "Flow detail records"). XaCCT treats network flow sources as asynchronous information sources that provide several of the essential UNIR parameters. However, there are parameters that would be expected from network flows but are currently lacking, parameters that may prove to be important for accounting and billing applications:

a) Network flow meters currently disregard the concept of "service" or "application" (provided for example by Checkpoint Technology's stateful-inspection accounting module) and only provide interfaces on the packet source and destination ports. Applications may utilize several ports during one session.

b) Network flows do not always provide a distinctive start/end-of-session time-stamp, so it may be difficult to obtain the network session duration.

XaCCT sees Network flow technology as a positive step in the endorsement of network session information provision by the industry, we look forward for the time when networking devices and applications shall provide CDR/SMDR-like records as a logging requirement standard.

Real-time UNIR Processing

XaCCT collection system enables the creation of "virtual" information source modules. These may be used to enhance UNIRs without necessarily probing an real information source. Instead, it can be used to associate auxiliary static information to UNIRs.

For example, association of UNIRs to geographic or entity locales may include performing IP/Netmask matching operations which are more naturally performed at program level rather then at the database. Another example is the association of price or price model to transactions. In modern telephony billing and customer service systems, it is common use to associate prices in "near real-time" to CDRs and store them with the price in the database. The reason is that if a customer calls and wants to get a current balance, the system will not be able to perform such a complex query on-line, on a database with several hundred million records. (In the future, application servers may provide a price model information record as part of their accounting logs [4]).

Virtual information source modules will enable real-time alerting based on network events, activation of external applications such as pagers, conversion of UNIRs to other formats for delivery in real-time to other systems such as network management and 3rd party accounting, security and log collection systems.

Application Road-map

Internet Auditing and Reporting

XaCCT Internet Auditing & Reporting is an application that assists in customizing and generating accounting reports and statistics based on the UNIRs stored in the database.. The system is accessible by the user through any Java enabled Web browser.

The system is supplied with a set of predefined queries and reports. Many users may find it sufficient to use these reports, rather than designing and generating their own reports. Should the user desire to customize, or add to these reports, this may be done using powerful tools for designing network-information-aware queries and reporting tools, supplied with the product. These tools are complemented by a set of business graphics to graphically render the query results (e.g. pie charts). The user is guided through creating complex queries and reports based on CEM generated tables, other tables in the database, mathematical and logical functions.

Reports may be produced either on-line, directly from the Web browser, or as batch reports, scheduled and recurring reports (every Day, Week, Month, etc.). Reports may be generated either as HTML documents or as delimited text files (suitable for import into various office automation tools). Reports may also be sent to users via E-mail or printed automatically to be used as invoices.

The queries and reports provide facilities to perform data-mining, to rapidly explore the detailed transactions of particular real-world entities. The system can be used to produce reports for modeling different business models to simulate payment scenarios over past periods in order to lower the risks in moving to new billing models. The format of the Network Session information in the database is available and can be accessed also by other 3rd party reporting facilities.

Future releases of the application will enable the activation of external applications when encountering UNIRs that match given expressions. This can be used to verify that the organization's security policy and Quality-of-Service policy has really been implemented. An alert in the form of e-mail, pager or 3rd party security system can be activated upon suspicious activity. Because of the unique technology of pipelined data collection (network session enhancement process), alerting can be achieved not only on raw network related information, but on higher level of information such as "category" of destination address, suspicious content of an e-mail or lack of Quality-of-Service.

Network Capacity Planning (NCP)

Using the unique view of the Intranet traffic provided by XaCCT's UNIR collection system, network managers will be able to plan and budget new network hardware and software, redistribute application-servers load and create periodic baselines of network consumption in order to view and predict trends, changes and exceptions in the network. The XaCCT NCP facilitates the creation of management reports for current status, trends and

what-if's scenarios in the WAN, extending XaCCT Internet Auditing & Reporting to accommodate for I.T. managers need for trend analysis, capacity planning and prediction.

Using XaCCT NCP, I.T. managers will be able to list for example, the top consuming Intranet bandwidth applications and users together with the effective bandwidth and latency actually experienced together with the average network load at the time. In a normal "best-effort" packet multiplexing technology, this information can be used to determine, weather the current ISP is [Internet wise] geographically close to the corporate's popular destination sites, and if not, recommend to move to another ISP that will provide better bandwidth or Quality-of-Service (bandwidth assurance) to the sites of interest. In a networking environment providing "assured bandwidth" through Quality-of-Service technologies, this information can provide proof of the received service. Such information can also be used to determine weather an application server is posing a bottle-neck and needs to be upgraded or should the server be distributed in order to lower the delay in service. It is not always clear why a client application seems "slow". By looking at trends of the application's effective throughput, it is easier to determine whether the bottleneck is with the network with the software.

In order to introduce Computer Telephony Integrated (CTI) and Video Conferencing systems or any new bandwidth consuming applications, to the corporate, academic or ISP networks, it is required to predict weather the current networking infrastructure can support the extra load and if not, what part of the network needs upgrading.

One possible method of predicting network future consumption and its behavioral trends, the XaCCT NCP may introduce a unique method taken after stock-quote prediction systems based on neural-networks. A neural net is trained with historic data composed of either raw UNIRs or aggregate and statistics tables. The neural network formulates a model or hypothesis of the UNIR parameters under analysis, the hypothesis being: future traffic data is based on present and past traffic data. The model is tested using historical data. For each day in a given history, the neural network makes a prediction based on information prior to that day, then the prediction is compared to what actually happen in the following twenty days or so of the date under test. If the prediction was correct then the model is reinforced, if it wasn't then the model is corrected. After enough iterations, the neural net is trained and the resulting model is a very refined representation of the network's behavior. Using this neural net, XaCCT NCP predicts future values of the given UNIR parameters based on up-to-date traffic parameters.

Another feature that will be supported by XaCCT NCP is generating views of real UNIRs taken from historical traffic data mixed with external UNIRs taken from some other system. The external UNIRs can be simulated and can provide a means of viewing the network's behavior when introducing extra network load. For example, one method of predicting bandwidth allocation due to CTI, is to merge historical Call Data Records (CDR) taken from the corporate PBX switches, into existing XaCCT collected network sessions. (The time and user ID are unique throughout the two systems so merging the data can be done based on those keys.). Another example is to replicate people's network usage records in order to predict network behavior when introducing new employees.

In the above methods of predicting future traffic and of merging historical traffic records with replicated data or external usage records the merged view can be used as input to usage and trend reports and data-mining applications.

The system comes with built-in support for business visualization graphics to assist in creating meaningful reports. The reports can also be exported to external office automation tools and legacy systems.

Accounting and Billing

The XaCCT Accounting and Billing application is a full scale billing application tailored to model various economic models for network usage. The application takes UNIRs collected through the XaCCT collection system and associates them with prices and billable entities.

Environment - The issuer of the bill has to start by defining various parameters related to the business environment such as currency, taxes, type of payment terms (such as credit card charging, cheque, money order, cache etc.), country specific parameters, and schedules for billing. There may be a need to communicate with local legacy systems such as SAP and Oracle Applications or with other centralized naming services such as X500 or LDAP from which users, customers and accounting information may be imported. In an ISP environment, the system interacts with RADIUS and TACACS(+) servers for all naming operations.

Billing mechanism for up to 400 bills is usually done by creating PDF files that are printed and sent to the customer. For a larger customer base systems, the billing parameters are exported to EDIFAC or GENCOD and sent to legacy billing systems or external billing systems. Users can receive their balance status through a web-browser (using industry standard encryption systems such as Netscape Server/Client RSA support). Users can receive their bills through an electronically signed e-mail attachment.

Tools will be provided for producing business statistics and reporting about general operation of the billing system, a sales commission system to monitor business unit performance, a settlement system to handle disputations with external network connected entities such as carriers, over lost data, lost sessions and failure to meet Quality-of-Service.

A customer service and support system will enable the operator to easily add or modify customers parameters, to provide on-line data as to customer's offerings and balance, to manually modify future bills following demands from customers concerning lost connections or bad QoS, and many other service features that may easily tailored to meet customer needs.

The systems also supports building what-if scenarios to test what would be the best pricing policy under given conditions and usage patterns. Special AI algorithms will be used to find the best/most profitable policies given ranges and limitations, again in order to assist in finding the best policy to suit both customer and organization and enable the organization to collate their profit with the real prices of deploying large WANs.

The following elements compose the basic structure and operation of XaCCT Billing system :

Policy - Policies are query rules (where clauses) that generate a subset (or a view) of the database of Network Sessions. Policies are the building blocks of the product elements offered to the customers. For example, a policy can generate a subset of the sessions that apply to the rule: "high-bandwidth sessions performed between 01:00PM and 08:00AM". This Policy can be charged with a special tariff since it may makes sense to charge less at night when there are less users logged in. The users is guided through a user friendly query builder for generating complex Policies.

Formula - A Formula calculates the price for a block of Network Sessions using various scalar and vector mathematical operators. For example, a formula named "Special reduction for hi-bandwidth usage at night" can multiply the total number of bytes of the sessions that apply to the previous example Policy, by some pre-calculated price/bandwidth ratio and multiply the result by 0.8 to give a 20% discount.

Product - A Product is an ordered set of Policies and for each Policy, an ordered set of Formulas. If a formula is inapplicable for a given block of network sessions, then the next Formula is taken, until the last Formula is is taken, which should be a default formula that applies to all cases. A Formula can be inapplicable to a set of Network Sessions if it depends on records that do not appear in the block. The Policies are an ordered set since a Policy may not necessarily apply to any group of network sessions as some of the records that the Policy depends on, may be missing, in which case the Policy will be inapplicable, so the next Policy will be tested against the network session block. This procedure continues until either an applicable Policy is found or the last catch-all Policy is reached.

Consumer - A Consumer is a person, computer, application, profit-center or any other uniquely defined entity producing Network Sessions. The Consumer ID is always set for every Network Transaction in the collection process. Each Consumer is associated to a Product and to a Customer (multiple users can belong to the same client company or department).

Customer - A customer is defined by a unique key, and parameters defining the contract, including a set of parameters such as payment terms, credit, address, billing contact, time and method etc.

Several features distinguish between the corporate and ISP accounting & billing systems and tailoring is expected to be necessary in most cases. XaCCT provides an unprecedented technology and business proposition to both worlds in terms of features and versatility, giving the network manager the power to view the data passed in the WAN from a fine-grained resolution granularity of small network sessions up to brusque statistics, and to be able to associate costs in easy user-friendly ways to any piece of information gathered about the network usage by users and applications.

Session Reconstruction module

First releases of XaCCT collection system will provide a unique technology for reconstructing network sessions from various sources of packet streams. The module will reside in possibly dedicated Gatherers and receive the flows by listening to the LAN, redirecting packets from routers and Firewalls (where applicable), and tapping onto WAN interfaces attached to routers using specially designed hardware interfaces.

The records gathered by the session reconstruction module will contain all the network related parameters that may be retrieved from a network session including session content.

Session content will enrich XaCCT applications by providing the ability to filter and aggregate sessions based on the content of the session.

In addition, UNIRs containing session content will introduce a new set of security related features that can record messages of interest, alert when noticing certain message contents appear, modify Firewall and Quality-of-Service policies when a certain event is detected.

Telco ISP Bridge

Glossary

CEM

The Central Event Manager (CEM) orchestrates the Gatherers thus controlling the information source configuration procedure and the UNIR enhancement processes. The CEM is connected to a database for storing UNIRs and configuration parameters.

Gatherer

Gatherers are light-weight smart-agent processes in non-dedicated computers. They probe an abundant number of information sources for Network Sessions, directory name services and other databases, transforming this information into UNIRs. Information collected is consolidated through an enhancement process associating the UNIRs to world entities.

Information Source

Almost all network devices and application servers provide some kind of logging or statistical information concerning their activity. An information source can be the log file of a Mail server, the logging facility of a Firewall, a traffic statistic available on a Router, a database entry accessible through the Web, captured packet data flow etc.

Synchronous IS are probed each time with different parameters and provide records of information in return. Asynchronous I.S. are set once and continuously deliver records of information.

Information Source Module

The Gatherer relies on information source modules to access the information source and is thus designed to be completely independent of the particular information sources that it interfaces with. information source modules are automatically upgraded from a central repository in the CEM.

Network Session

A "Network Session" is the set of parameters defining the data that are passed back and forth between computers when a client-server application is activated. These parameters include source and destination IP addresses, application type, related people or profit & loss entities, duration, information about the content, etc.

Quality of Service (QoS)

A term used in the networking industry for methods of assuring end-to-end bandwidth for a predefined network session. This technology is essential for business CTI applications.

UNIR - Unified Network Information Record

The UNIR is a set of "ideal" parameters that may or may not be provided by a given information source. A TACACS authentication server provides user names and associated IP addresses but does not provide Network Session information. IOS 10.X CISCO network-statistics provide traffic matrix bandwidth consumption over a period of time, but do not associate an informative name to the source and destination addresses.

UNIR Enhancement

The procedure by which XaCCT acquires information from several information sources and merges them to produce more meaningful UNIRs. The enhancement is a process of channeling the UNIRs from an asynchronous streaming information source (such as a router or a firewall), through other Gatheres that access synchronous information sources (such as databases) and add additional information to the records that were previously accumulated. The final stop in the chain is the CEM who stores the enhanced UNIRs in the database. The enhancement process is highly optimized and is controlled by the CEM.

For example, in an enhancement procedure, the UNIRs taken from a CISCO traffic matrix (source IP, destination IP, nbytes, npackets) can be enhanced by transforming the source IP into a user name by contacting an LDAP naming server. The destination IP can be converted to its DNS equivalent and can then be associated with a company name by querying the WHOIS database.

References

[1] Brownlee, N., " Traffic Flow Measurement: Experiences with NeTraMet", RFC 2123, The University of Auckland, March 1997.

[2] Brownlee, N., Mills, C., and G. Ruth, "Traffic Flow Measurement: Architecture", RFC 2063, The University of Auckland, Bolt Beranek and Newman Inc., GTE Laboratories, Inc, January 1997.

[3] Cisco Systems, "NetFlow: White Paper", Cisco Systems Inc, 1997

[4] Msix Initiative, "http://www.m6.org"